

# (Quasi)Periodicity Quantification in Video Data, Using Topology

Christopher J. Tralie<sup>†</sup> and Jose A. Perea<sup>\*</sup>

**Abstract**—This work introduces a novel framework for quantifying the presence and strength of recurrent dynamics in video data. Specifically, we provide continuous measures of periodicity (perfect repetition) and quasiperiodicity (superposition of periodic modes with non-commensurate periods), in a way which does not require segmentation, training, object tracking or 1-dimensional surrogate signals. Our methodology operates directly on video data. The approach combines ideas from nonlinear time series analysis (delay embeddings) and computational topology (persistent homology), by translating the problem of finding recurrent dynamics in video data, into the problem of determining the circularity or toroidality of an associated geometric space. Through extensive testing, we show the robustness of our scores with respect to several noise models/levels; we show that our periodicity score is superior to other methods when compared to human-generated periodicity rankings; and furthermore, we show that our quasiperiodicity score clearly indicates the presence of biphonation in videos of vibrating vocal folds.

## I. INTRODUCTION

Periodicity characterizes many natural motions including animal locomotion (walking/wing flapping/slithering), spinning wheels, oscillating pendulums, etc. Quasiperiodicity, or the presence of non-commensurate frequencies, occurs naturally during transitions from ordinary to chaotic dynamics [11]. In this work, we automate the analysis of videos capturing periodic and quasiperiodic motion. To identify both classes of motion in a unified framework, we generalize 1-dimensional (1D) sliding window embeddings [38] to reconstruct periodic and quasiperiodic attractors from videos. We analyze the resulting attractors using persistent homology, a technique which combines geometry and topology (Section II-B), and we return a single number in the range  $[0, 1]$  that indicates the degree of periodicity or quasiperiodicity in the corresponding video.

Many approaches to periodicity analysis derive 1D signals from objects tracked in the video. By contrast, our technique operates on raw pixels, sidestepping tracking; as such, we propose a general and domain independent approach. We show that our periodicity measure compares favorably to others in the literature when ranking videos (Section IV-B). Furthermore, to our knowledge, there is no other method able to quantify the existence of quasiperiodicity directly from video data, and we exploit this to measure the degree of “biphonation” in

videos of vibrating vocal folds (Section IV-C), which is useful in automatically diagnosing speech pathologies. In the context of applied topology, our quasiperiodicity score is one of the first applications of persistent  $H_2$  to high dimensional data, which is largely possible due to recent advancements in the computational feasibility of persistent homology [3].

### A. Prior Work on Recurrence in Videos

1) *1D Surrogate Signals*: One common strategy is to derive a 1D function from the video to act as a surrogate for its dynamics, and then to use either frequency domain (Fourier transform) or time domain (autocorrelation, peak finding) techniques. One of the earliest works in this genre finds level set surfaces in a spatiotemporal “XYT” volume of video (all frames stacked on top of each other), and then uses curvature scale space on curves that live on these “spatiotemporal surfaces” as the 1D function [1]. [34] use Fourier Transforms on pixels which exhibit motion, and define a measure of periodicity based on the energy around the Fourier peak and its harmonics. [10] extract contours and find eigenshapes from the contours to classify and parameterize motion within a period. Frequency estimation is done by using Fourier analysis and peak detection on top of other 1D statistics derived from the contours, such as area and center of mass. Finally, [47] derive a 1D surrogate function based on mutual information between the first and subsequent frames, and then look for peaks in the similarity function with the help of a watershed method.

2) *Self-Similarity Matrices*: Another class of techniques relies on analyzing self-similarity matrices (SSMs) between frames, where similarity can be defined in a variety of ways. [37] track a set of points on a foreground object and compare them with an affine invariant similarity. Another widely recognized technique for periodicity quantification [6], derives periodicity measures based on self-similarity matrices of L1 pixel differences. This technique has inspired a diverse array of applications, including analyzing the cycles of expanding/contracting jellyfish [33], analyzing bat wings [2], and analyzing videos of autistic spectrum children performing characteristic repetitive motions such as “hand flapping” [22]. We compare to this technique in Section IV-B.

3) *Miscellaneous Techniques for Periodic Video Quantification*: There are also a number of works that don’t fall into the two categories above. Some works focus solely on walking humans, since that is one of the most common types of periodic motion in videos of interest to people. [29] look at the “braiding patterns” that occur in XYT slices of videos of walking people. [17] perform blob tracking on the foreground

<sup>†</sup> Department of Electrical And Computer Engineering, Duke University, Durham, NC, USA. e-mail: ctrialie@alumni.princeton.edu

<sup>\*</sup> Department of Mathematics and Department of Computational Mathematics, Science & Engineering, Michigan State University, East Lansing, MI, USA. e-mail: joperea@math.msu.edu

of a walking person, and use the ratio of the second and first eigenvalues of PCA on that blob.

For more general periodic videos, [43] make a codebook of visual words and look for repetitions within the resulting string. [23] take a deep learning approach to counting the number of periods that occur in a video segment. They use a 3D convolutional neural network on spatially downsampled, non-sequential regions of interest, which are uniformly spaced in time, to estimate the length of the cycle. Finally, perhaps the most philosophically similar work to ours is the work of [41], who use cohomology to find maps of MOCAP data to the circle for parameterizing periodic motions, though this work does not provide a way to quantify periodicity.

4) *Our Work*: Geometry provides a natural way to quantify periodicity by measuring the roundness of delay embeddings. Crucially, it also provides way to measure *quasiperiodicity*, for which quantitative approaches are lacking. As we show in Section IV-C, this simplifies the pipeline for detecting anomalies in videos of vibrating vocal folds. Finally, in contrast to both frequency and time domain techniques, our method does not rely on the period length being an integer multiple of the sampling rate.

## II. BACKGROUND

### A. Delay Embeddings And Their Geometry

We are now ready to present our construction.

1) *Video Delay Embeddings*: We will regard a video as a sequence of grayscale<sup>1</sup> image frames indexed by the positive real numbers. That is, given positive integers  $W$  (width) and  $H$  (height), a video with  $W \times H$  pixels is a function

$$X : \mathbb{R}^+ \longrightarrow \mathbb{R}^{W \times H}$$

In particular, a sequence of images  $X_1, X_2, \dots \in \mathbb{R}^{W \times H}$  sampled at discrete times  $t_1 < t_2 < \dots$  yields one such function via interpolation. For an integer  $d \geq 0$ , known as the *dimension*, a real number  $\tau > 0$ , known as the *delay*, and a video  $X : \mathbb{R}^+ \longrightarrow \mathbb{R}^{W \times H}$ , we define the **sliding window** (also referred to as time delay) **embedding** of  $X$  – with parameters  $d$  and  $\tau$  – at time  $t \in \mathbb{R}^+$  as the vector

$$SW_{d,\tau}X(t) = \begin{bmatrix} X(t) \\ X(t+\tau) \\ \vdots \\ X(t+d\tau) \end{bmatrix} \in \mathbb{R}^{W \times H \times (d+1)} \quad (1)$$

The subset of  $\mathbb{R}^{W \times H \times (d+1)}$  resulting from varying  $t$  will be referred to as the sliding window embedding of  $X$ . We remark that since the pixel measurement locations are fixed, the sliding window embedding is an “Eulerian” view into the dynamics of the video. Note that delay embeddings are generally applied to 1D time series, which can be viewed as 1-pixel videos ( $W = H = 1$ ) in our framework. Hence equation (1) is essentially the concatenation of the delay embeddings of each individual

pixel in the video into one large vector. One of the main points we leverage in this paper is the fact that the geometry of the sliding window embedding carries fundamental information about the original video. We explore this next.

2) *Geometry of 1-Pixel Video Delay Embeddings*: In order to describe the geometry of sliding window embeddings of videos with arbitrary resolution  $W \times H$ , we start with two simple examples in the case of 1D time series (1-pixel videos). Consider the *harmonic* (i.e. periodic) signal

$$f_h(t) = \cos\left(\frac{\pi}{5}t\right) + \cos\left(\frac{\pi}{15}t\right) \quad (2)$$

and the *quasiperiodic* signal

$$f_q(t) = \cos\left(\frac{\pi}{5}t\right) + \cos\left(\frac{1}{5}t\right) \quad (3)$$

The signal  $f_h$  is referred to as harmonic because its constitutive frequencies,  $\frac{1}{10}$  and  $\frac{1}{30}$ , are *commensurate*; that is, they are linearly dependent over the rational numbers  $\mathbb{Q} \subset \mathbb{R}$ . By way of contrast, the underlying frequencies of the signal  $f_q$ ,  $\frac{1}{10}$  and  $\frac{1}{10\pi}$ , are linearly independent over  $\mathbb{Q}$  and hence *non-commensurate*. We use the term quasiperiodicity, as in the non-linear dynamics literature [19], to denote the superposition of periodic processes whose frequencies are non-commensurate. This differs from other definitions in the literature (e.g. [43]) which regard quasiperiodic as any deviation from perfect repetition.

A geometric argument from [31] (see equation 7 below and the discussion that follows) shows that given a periodic function  $f : [0, 2\pi] \longrightarrow \mathbb{R}$  with exactly  $N$  harmonics, if  $d \geq 2N$  and  $0 < \tau < \frac{2\pi}{d}$  then the sliding window embedding  $SW_{d,\tau}f$  is a topological circle (i.e. a closed curve without self-intersections) which wraps around an  $N$ -dimensional torus

$$\mathbb{T}^N = \underbrace{S^1 \times \dots \times S^1}_{N\text{-times}}, \quad S^1 = \{z \in \mathbb{C} : |z| = 1\}$$

As an illustration, we show in Figure 1 a plot of  $f_h$  and of its sliding window embedding  $SW_{d,\tau}f_h$ , via a PCA (Principal Component Analysis) 3-dimensional projection.




Fig. 1: Sliding window embedding of the harmonic signal  $f_h$ . Colors in the signal correspond to colors of the points in the PCA plot. The sliding window embedding traces a topological circle wrapped around a 2-dimensional torus.

However, if  $g : \mathbb{R} \longrightarrow \mathbb{R}$  is quasiperiodic with  $N$  distinct non-commensurate frequencies then, for appropriate  $d$  and  $\tau$ ,  $SW_{d,\tau}g$  is dense in (i.e. fills out)  $\mathbb{T}^N$  [30]. Figure 2 shows a

<sup>1</sup>For color videos we can treat each channel independently, yielding a vector in  $\mathbb{R}^{W \times H \times 3}$ . In practice, there isn't much of a difference between color and grayscale embeddings in our framework for the videos we consider.

plot of the quasiperiodic signal  $f_q(t)$  and a 3-dimensional projection, via PCA, of its sliding window embedding  $SW_{d,\tau}f_q$ .

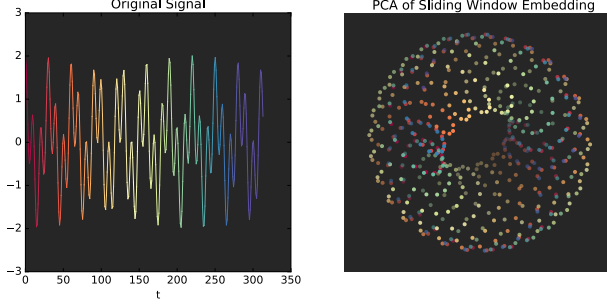


Fig. 2: Sliding window embedding of the quasiperiodic signal  $f_q$ . Colors in the signal correspond to colors of the points in the PCA plot. The sliding window embedding is dense in a 2-dimensional torus.

The difference in geometry of the delay embeddings is stark compared to the difference between their power spectral densities, as shown in Figure 3.




Fig. 3: The power spectral densities of 300 sample commensurate and non-commensurate signals with relative harmonics at ratios 3 and  $\pi$ , respectively. The difference is not nearly as evident as it is with the geometry of their sliding window embeddings. Additionally, unless sampling is commensurate with a frequency, a fixed Fourier basis causes that frequency component to bleed into many frequency bins in a sinc-like pattern, making precise peak finding difficult.

Moreover, as we will see next, the interpretation of periodicity and quasiperiodicity as circularity and toroidality of sliding window embeddings remains true for videos with higher resolution (i.e.  $\max\{W, H\} > 1$ ). The rest of the paper will show how one can use *persistent homology*, a tool from the field of computational topology, to quantify the presence of (quasi)periodicity in a video by measuring the geometry of its associated sliding window embedding. In short, we propose a *periodicity score* which measures the degree to which a collection of points spans a topological circle, and a *quasiperiodicity score* which quantifies the degree to which a set of points covers a torus. This approach will be validated extensively: we show that our (quasi)periodicity detection method is robust under several noise models (motion blur, additive Gaussian white noise, and MPEG bit corruption); we compare several periodicity quantification algorithms and

show that our approach is the most closely aligned with human subjects; finally, we provide an application to the automatic classification of dynamic regimes in high-speed laryngeal video-endoscopy.

3) *Geometry of Video Delay Embeddings*: Though it may seem daunting compared to the 1D case, the geometry of the delay embedding shares many similarities for periodic videos, as shown in [39]. Let us argue why sliding window embeddings from (quasi)periodic videos have the geometry we have described so far. To this end, consider an example video  $X$  that contains a set of  $N$  frequencies  $\omega_1, \omega_2, \dots, \omega_N$ . Let the amplitude of the  $n^{\text{th}}$  frequency and  $i^{\text{th}}$  pixel be  $a_{in}$ . For simplicity, assume that each is a cosine with a zero phase offset (this simplifies the math without loss of generality). Then the time series at pixel  $i$  can be written as

$$X_i(t) = \sum_{n=1}^N a_{in} \cos(\omega_n t) \quad (4)$$

Grouping all of the coefficients together into a  $(W \times H) \times N$  matrix  $A$ , we can write

$$X(t) = \sum_{n=1}^N A^n \cos(\omega_n t) \quad (5)$$

where  $A^n$  stands for the  $n^{\text{th}}$  column of  $A$ . Now, we construct a delay embedding as in Equation 1:

$$SW_{d,\tau}X(t) = \sum_{n=1}^N \begin{bmatrix} A^n \cos(\omega_n t) \\ \vdots \\ A^n \cos(\omega_n(t + d\tau)) \end{bmatrix} \quad (6)$$

applying the cosine sum identity, we get

$$SW_{d,\tau}X(t) = \sum_{n=1}^N \vec{u}_n \cos(\omega_n t) - \vec{v}_n \sin(\omega_n t) \quad (7)$$

where  $\vec{u}_n, \vec{v}_n \in \mathbb{R}^{W \times H \times (d+1)}$  are constant vectors. In other words, the sliding window embedding of this video is the sum of linearly independent ellipses, which lie in the space of  $d+1$  frame videos at resolution  $W \times H$ . As shown in [31] for the case of commensurate frequencies, when the window length is just under the length of the period, all of the  $\vec{u}_n$  and  $\vec{v}_n$  vectors become orthogonal, and so they can be recovered by doing PCA on  $SW_{d,\tau}X(t)$ . Figure 4 shows the components of the first 8 PCA vectors for a horizontal line of pixels in a video of an oscillating pendulum. Note how the oscillations are present both temporally and spatially.

4) *The High Dimensional Geometry of Repeated Pulses*: Using Eulerian coordinates has an important impact on the geometry of delay embeddings of natural videos. As Figure 5 shows, pixels often jump from foreground to background in a pattern similar to square waves.

These types of abrupt transitions require higher dimensional embeddings to reconstruct the geometry. To see why, first extract one period of a signal with period  $\ell$  at a pixel  $X_i(t)$ :

$$f_i(t) = \begin{cases} X_i(t) & 0 \leq t \leq \ell \\ 0 & \text{otherwise} \end{cases} \quad (8)$$




Fig. 4: Showing an XT slice of the principal components on  $SW_{d,1}$  for the synthetic video of an oscillating pendulum, where  $d$  is chosen just under the length of the period of oscillation ( $\sim 25$  frames).

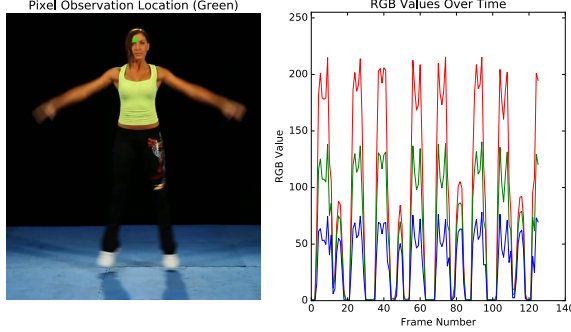


Fig. 5: An example of an Eulerian pixel witnessing a foreground/background transition in a video of a woman doing jumping jacks. Red, green, and blue channels are plotted over time. These transitions induce a per pixel periodic signal with sharp transitions, which leads to high dimensionality in an appropriate sliding window embedding.

Then  $X_i(t)$  can be rewritten in terms of the pulse as

$$X_i(t) = \sum_{m=-\infty}^{\infty} f_i(t - m\ell) \quad (9)$$

Since  $X_i(t)$  repeats itself, regardless of what  $f_i(t)$  looks like, *periodic summation* discretizes the frequency domain [32]

$$\mathcal{F}\{X_i(t)\}(k) \propto \sum_{m=-\infty}^{\infty} \mathcal{F}(f_i(t)) \left(\frac{m}{\ell}\right) \delta\left(\frac{m}{\ell} - k\right) \quad (10)$$

Switching back to the time domain, we can write  $X_i(t)$  as

$$X_i(t) \propto \sum_{m=-\infty}^{\infty} \mathcal{F}(f_i(t)) \left(\frac{m}{\ell}\right) e^{i\frac{2\pi m}{\ell}t} \quad (11)$$

In other words, each pixel is the sum of some constant offset plus a (possibly infinite) set of harmonics at integer multiples of  $\frac{1}{\ell}$ . For instance, applying Equation 11 to a square wave of period  $\ell$  centered at the origin is a roundabout way of deriving the Fourier Series

$$\sin\left(\frac{2\pi}{\ell}t\right) + \frac{1}{3}\sin\left(\frac{6\pi}{\ell}t\right) + \frac{1}{5}\sin\left(\frac{10\pi}{\ell}t\right) + \dots \quad (12)$$

by sampling the sinc function  $\sin(\pi\ell f)/(\pi f)$  at intervals of  $m/2\ell$  (every odd  $m$  coincides with  $\pi/2 + k\pi$ , proportional to  $1/k$ , and every even harmonic is zero coinciding with  $\pi k$ ). In general, the sharper the transitions are in  $X_i(t)$ ,

the longer the tail of  $\mathcal{F}\{f_i(t)\}$  will be, and the more high frequency harmonics will exist in the embedding, calling for a higher delay dimension to fully capture the geometry, since every harmonic lives on a linearly independent ellipse. Similar observations about harmonics have been made in images for collections of patches around sharp edges ([48], Figure 2).

## B. Persistent Homology

Informally, topology is the study of properties of spaces which do not change after stretching without gluing or tearing. For instance, the number of connected components and the number of (essentially different) 1-dimensional loops which do not bound a 2-dimensional disk, are both topological properties of a space. It follows that a circle and a square are topologically equivalent since one can deform one onto the other, but a circle and a line segment are not because that would require either gluing the endpoints of the line segment or tearing the circle. Homology [12] is a tool from algebraic topology designed to measure these types of properties, and persistent homology [50] is an adaptation of these ideas to discrete collections of points (e.g., sliding window embeddings). We briefly introduce these concepts next.

1) *Simplicial Complexes*: A simplicial complex is a combinatorial object used to represent and discretize a continuous space. With a discretization available, one can then compute topological properties by algorithmic means. Formally, a **simplicial complex** with vertices in a nonempty set  $V$  is a collection  $K$  of nonempty finite subsets  $\sigma \subset V$  so that  $\emptyset \neq \tau \subset \sigma \in K$  always implies  $\tau \in K$ . An element  $\sigma \in K$  is called a **simplex**, and if  $\sigma$  has  $(n+1)$  elements then it is called an  $n$ -simplex. The cases  $n = 0, 1, 2$  are special, 0-simplices are called vertices, 1-simplices are called edges and 2-simplices are called faces. Here is an example to keep in mind: the circle  $S^1 = \{z \in \mathbb{C} : |z| = 1\}$  is a continuous space but its topology can be captured by a simplicial complex  $K$  with three vertices  $a, b, c$ , and three edges  $\{a, b\}, \{b, c\}, \{a, c\}$ . That is, in terms of topological properties, the simplicial complex

$$K = \{\{a\}, \{b\}, \{c\}, \{a, b\}, \{b, c\}, \{a, c\}\}$$

can be regarded as a combinatorial surrogate for  $S^1$ .

2) *Simplicial Homology*: The homology of a simplicial complex provides a measure for its topological structure: e.g., number of connected components, existence of 1-dimensional cycles which are not filled in by a 2-dimensional surface, closed 2-dimensional surfaces bounding a 3-dimensional void, etc. We define this concept next. Given a simplicial complex  $K$  fix an integer  $n \geq 0$  and a field  $\mathbb{F}$  (e.g.,  $\mathbb{Q}, \mathbb{R}, \mathbb{C}$  or  $\mathbb{Z}_p = \mathbb{Z}/p\mathbb{Z}$  for  $p \in \mathbb{N}$  prime). An  $n$ -chain of  $K$  with coefficients in  $\mathbb{F}$  is a formal sum

$$\sum_{i=1}^I a_i \sigma_i, \quad I \in \mathbb{N}$$

where each  $\sigma_i$  is an  $n$ -simplex of  $K$  and  $a_i \in \mathbb{F}$ . The set of  $n$ -chains is a vector space over  $\mathbb{F}$ , which we denote by  $C_n(K; \mathbb{F})$  and refer to as the space of  $n$ -chains. Notice that

the  $n$ -simplices form a basis for  $C_n(K; \mathbb{F})$ . The boundary of an  $n$ -simplex  $\sigma = \{u_0, \dots, u_n\}$  is the  $(n-1)$ -chain

$$\partial_n \sigma = \sum_{j=0}^n (-1)^j \{u_0, u_1, \dots, \hat{u}_j, \dots, u_n\} \quad (13)$$

where  $\hat{u}_j$  indicates the omission of the vertex  $u_j$ . Extending  $\partial_n$  linearly to  $n$ -chains yields a well-defined linear transformation  $\partial_n : C_n(K; \mathbb{F}) \rightarrow C_{n-1}(K; \mathbb{F})$  called the  $n$ -boundary map and which, as one can check, satisfies  $\partial_{n-1} \circ \partial_n = 0$  for all  $n > 0$ . Hence, if  $B_n(K; \mathbb{F}) = \text{Im}(\partial_{n+1})$  denotes the space of  $n$ -boundaries and  $Z_n(K; \mathbb{F}) = \text{Ker}(\partial_n)$  the space of  $n$ -cycles then  $B_n(K; \mathbb{F}) \subset Z_n(K; \mathbb{F})$ . The  $n$ -th **homology** of the simplicial complex  $K$ , with coefficients in the field  $\mathbb{F}$ , is defined as the quotient of these two vector spaces:

$$H_n(K; \mathbb{F}) = Z_n(K; \mathbb{F}) / B_n(K; \mathbb{F})$$

which is itself a vector space over  $\mathbb{F}$ . Elements of  $H_n(K; \mathbb{F})$  are referred to as homology classes, and two  $n$ -cycles  $\alpha, \alpha'$  yield distinct homology classes if and only if  $\alpha - \alpha'$  is not an  $n$ -boundary. Intuitively,  $\dim(H_0(K; \mathbb{F}))$  counts the number of connected components in  $K$ ,  $\dim(H_1(K; \mathbb{F}))$  counts the number of independent (i.e. non-homologous) holes, and  $\dim(H_2(K; \mathbb{F}))$  counts the number of independent voids. Going back to the simplicial complex  $K = \{\{a\}, \{b\}, \{c\}, \{a, b\}, \{a, c\}, \{b, c\}\}$  which serves as a surrogate for the unit circle  $S^1$ , one can check that  $H_0(K; \mathbb{F}) \cong \mathbb{F}$ ,  $H_1(K; \mathbb{F}) \cong \mathbb{F}$  and that  $H_n(K; \mathbb{F}) = \{0\}$  for all  $n \geq 2$ . This is consistent with the topology of  $S^1$ : it is connected, has one hole in dimension one and no other features in higher dimensions. However, if  $K$  discretizes the 2-torus  $\mathbb{T}^2$  (e.g., the natural triangulation with 9 vertices, 27 edges and 18 faces) then  $H_0(K; \mathbb{F}) \cong \mathbb{F}$ ,  $H_1(K; \mathbb{F}) \cong \mathbb{F}^2$ ,  $H_2(K; \mathbb{F}) \cong \mathbb{F}$  and  $H_n(K; \mathbb{F}) = \{0\}$  for  $n \geq 3$ . This reflects the fact that  $\mathbb{T}^2$  is connected, it has two independent circular holes (in the vertical and horizontal directions, resp), bounds a 3-dimensional void and has no other features in higher dimensions. We remark that determining the homology of a (finite) simplicial complex reduces to a problem in linear algebra. Indeed,  $\dim(H_n(K; \mathbb{F})) = \text{null}(\partial_n) - \text{rank}(\partial_{n+1})$  and an explicit basis for  $H_n(K; \mathbb{F})$  can be computed via elementary row/column operations on the matrices of  $\partial_n$  and  $\partial_{n+1}$ .

3) *Persistent Homology of Point Clouds*: In practice the sliding window embedding of a video  $X$  is a finite set  $\text{SW}_{d,\tau} X = \{SW_{d,\tau} X(t) : t \in T\}$ , determined by a choice of  $T \subset \mathbb{R}$  finite. Moreover, since  $\text{SW}_{d,\tau} X \subset \mathbb{R}^{W \times H \times (d+1)}$  then the restriction of the ambient Euclidean distance endows  $\text{SW}_{d,\tau} X$  with the structure of a finite metric space. Discrete metric spaces, also referred to as *point clouds*, are trivial from a homological point view: a point cloud with  $N$  points simply has  $N$  connected components and no other features (e.g., holes) in higher dimensions. However, when a point cloud has been sampled from/around a continuous space with non-trivial topology (e.g., a circle or a torus), one would expect that appropriate simplicial complexes with vertices on the point cloud should reflect the topology of the underlying continuous space. This is what we will exploit next. The idea is to build a multi-scale family of simplicial complexes, and then compute

homology across scales. In this setting, homology classes which persist through a wide range of scales are associated with significant topological features, while unstable homology classes are attributed to sampling artifacts.

Specifically, a **filtered simplicial complex** is an indexed family  $\mathcal{K} = \{K_\epsilon\}_{\epsilon \geq 0}$  of simplicial complexes so that  $\epsilon \leq \epsilon'$  always implies  $K_\epsilon \subset K_{\epsilon'}$ . Hence, if  $\epsilon \leq \epsilon'$  then

$$Z_n(K_\epsilon; \mathbb{F}) \subset Z_n(K_{\epsilon'}; \mathbb{F}) \quad \text{and} \quad B_n(K_\epsilon; \mathbb{F}) \subset B_n(K_{\epsilon'}; \mathbb{F})$$

Let  $\alpha \in Z_n(K_\epsilon; \mathbb{F})$  be so that  $\alpha \notin B_n(K_\epsilon; \mathbb{F})$ . We say that  $\alpha$  dies entering  $K_\epsilon$ ,  $\epsilon < \underline{\epsilon}$  if  $\alpha \in B_n(K_{\epsilon'}; \mathbb{F})$  occurs if and only if  $\epsilon' \geq \underline{\epsilon}$ . Similarly,  $\alpha \in Z_n(K_\epsilon; \mathbb{F})$  is said to have been born at  $K_\epsilon$  if there is no  $n$ -cycle  $\alpha' \in Z_n(K_{\epsilon'}; \mathbb{F})$ ,  $\epsilon' < \epsilon$ , so that  $\alpha - \alpha' \in B_n(K_\epsilon; \mathbb{F})$ . With this notation,  $\alpha$  is said to yield an  $n$ -dimensional **persistent homology** class for the filtered complex  $\mathcal{K}$ , with *birth-time*  $\epsilon$ , *death-time*  $\underline{\epsilon}$  and *persistence*  $\underline{\epsilon} - \epsilon$ . Moreover, two  $n$ -cycles  $\alpha$  and  $\alpha'$  are said to yield equivalent persistent homology classes if their birth and death times are the same,  $\epsilon$  and  $\underline{\epsilon}$  respectively, and there exists  $\epsilon < \epsilon' < \underline{\epsilon}$  so that the  $n$ -cycle  $\alpha - \alpha'$  dies entering  $K_{\epsilon'}$ . The collection of resulting birth-death pairs  $(\epsilon, \underline{\epsilon})$ , from all non-equivalent  $n$ -dimensional persistent homology classes, yields a multi-set  $\text{dgm}_n(\mathcal{K}; \mathbb{F}) \subset \mathbb{R}^2$  called the  $n$ -dimensional **persistence diagram** of  $\mathcal{K}$ . It is a multi-set in the sense that some of its elements may appear with repetitions (i.e. with multiplicity) as two non-equivalent persistent homology classes can have the same birth-death pair. Since a persistence diagram  $\text{dgm}$  is just a collection of points in the region  $\{(x, y) \in \mathbb{R}^2 : 0 \leq x < y\}$ , we will visualize it as a scatter plot. We will also include the diagonal  $y = x$  in order to visually convey the persistence  $d - b$  of each birth-death pair  $(b, d) \in \text{dgm}$ . In this setting, points far from the diagonal (i.e. with large persistence) represent topological features which are stable across scales and hence deemed significant, while points near the diagonal (i.e. with small persistence) are often associated with unstable features.

We remark that the computational task of determining all non-equivalent persistent homology classes of a filtered simplicial complex can, surprisingly, be reduced to computing the homology of a single simplicial complex [3], [50]. This, as we mentioned earlier, is a problem in linear algebra that can be solved via elementary row and column operations on appropriate boundary matrices.

Now back to our geometric setting. Given a point cloud  $(\mathbb{X}, d_{\mathbb{X}})$  – where  $\mathbb{X}$  is a finite set and  $d_{\mathbb{X}} : \mathbb{X} \times \mathbb{X} \rightarrow [0, \infty)$  is a distance function – the *Vietoris-Rips complex* (or Rips complex for short) at scale  $\epsilon \geq 0$  is the collection of non-empty subsets of  $\mathbb{X}$  with diameter less than or equal to  $\epsilon$ :

$$R_\epsilon(\mathbb{X}) := \{\sigma \subset \mathbb{X} : d_{\mathbb{X}}(x_1, x_2) \leq \epsilon, \forall x_i, x_j \in \sigma\} \quad (14)$$

Notice that  $R_\epsilon(\mathbb{X})$  is in fact a simplicial complex for all  $\epsilon \geq 0$  and that  $\mathcal{R}(\mathbb{X}) = \{R_\epsilon(\mathbb{X})\}_{\epsilon \geq 0}$  is a filtered simplicial complex, known as the *Rips filtration* of  $\mathbb{X}$ . We illustrate in Figure 6 the process of going from a point cloud to the 1-dimensional persistence diagram of its Rips filtration.

The persistent homology of the Rips filtration on  $\text{SW}_{d,\tau} X$ , and in particular its  $n$ -dimensional persistence diagrams for



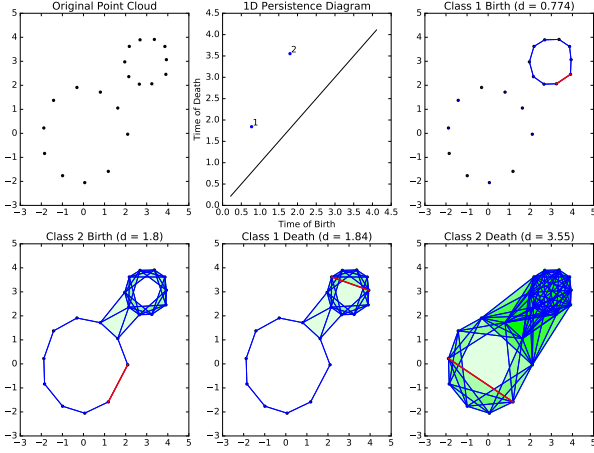


Fig. 6: From a point cloud to the 1-dimensional persistence diagram of its Rips filtration. Connected edges in the Rips filtration are drawn in blue, the birth/death of a class is indicated in red, and filled triangles are shaded green.

$n = 1, 2$ , are the objects we will use to quantify periodicity and quasiperiodicity in a video  $X$ . Figures 7 and 8 show the persistence diagrams of the Rips filtrations, on the sliding window embeddings, for the commensurate and non-commensurate signals from Figures 1 and 2, respectively. We use fast new code from the “Ripser” software package to make persistent  $H_2$  computation feasible [3].




Fig. 7: Sliding window embedding of the harmonic signal  $f_h$  (left) and the  $n$ -dimensional persistence diagrams  $n = 1, 2$  (right) of the associated Rips filtration. The sliding window embedding  $SW_{d,\tau}f_h$  traces a topological circle wrapped around a 2-dimensional torus. The persistence diagram in dimension one ( $H_1$ ) shows only one birth-death pair with prominent persistence; this is consistent with a point cloud sampled around a space with the topology of a circle.

### III. IMPLEMENTATION DETAILS

#### A. Reducing Memory Requirements with SVD

The first thing we will address is the memory requirements and ensuing computational burden to construct and access the sliding window embedding. As an illustration, suppose we have a video which has been discretely sampled at  $N$  different frames at a resolution of  $W \times H$ , and we do a delay embedding with dimension  $d$ , for some arbitrary  $\tau$ . Assuming 32 bit floats per grayscale value, storing the sliding window embedding




Fig. 8: Sliding window embedding of the quasiperiodic signal  $f_q$  (left) and the  $n$ -dimensional persistence diagrams  $n = 1, 2$  (right) of the associated Rips filtration. The sliding window embedding  $SW_{d,\tau}f_q$  is dense on a 2-dimensional torus. The persistence diagram in dimension one ( $H_1$ ) shows two birth-death pairs with prominent persistence, while the persistence diagram in dimension two ( $H_2$ ) shows one prominent birth-death pair; this is consistent with a point cloud sampled around a space with the topology of a 2-dimensional torus.

requires  $4WHN(d+1)$  bytes. For a low resolution  $200 \times 200$  video only 10 seconds long at 30fps, using  $d = 30$  already exceeds 1GB of memory.

Nonetheless, since constructing the Rips filtration only requires pairwise distances between different delay vectors, this enables a few optimizations. First of all, for  $N$  points in  $\mathbb{R}^{WH}$ , where  $N \ll WH$ , there exists an  $N$ -dimensional linear subspace which contains them. In particular let  $A$  be the  $(W \times H) \times N$  matrix with each video frame along a column. Performing a singular value decomposition  $A = USV^T$ , yields a matrix  $U$  whose columns form an orthonormal basis for the aforementioned  $N$ -dimensional linear subspace. Hence, by finding the coordinates of the original frame vectors with respect to this orthogonal basis

$$\hat{A} = U^T A = U^T USV = SV \quad (15)$$

and using the coordinates of the columns of  $SV$  instead of the original pixels, we get a sliding window embedding of lower dimension

$$SW_{d,\tau}(t) = \begin{bmatrix} U^T X(t) \\ \vdots \\ U^T X(t + d\tau) \end{bmatrix} \quad (16)$$

for which

$$\|SW_{d,\tau}(t) - SW_{d,\tau}(t')\| = \|SW_{d,\tau}X(t) - SW_{d,\tau}X(t')\|$$

Note that  $SV$  can be computed by finding the eigenvectors of  $A^T A$ ; this has a cost of  $O(W^2 H^2 + N^3)$  which is dominated by  $W^2 H^2$  if  $WH \gg N$ . In our example above, this alone reduces the memory requirements from 1GB to 10MB. Of course, this procedure is the most effective for short videos where there are actually many fewer frames than pixels, but this encompasses most of the examples in this work. In fact, the break even point for a  $200 \times 200$  30fps video is 22 minutes. A similar approach was used in the classical work on Eigenfaces [40] when computing the principal components over a set of face images.

### B. Distance Computation via Diagonal Convolutions

A different optimization is possible if  $\tau = 1$ ; that is, if delays are taken exactly on frames and no interpolation is needed. In this case, the squared Euclidean distance between  $SW_{d,1}(i)$  and  $SW_{d,1}(j)$  is

$$\|SW_{d,1}(i) - SW_{d,1}(j)\|_2^2 = \sum_{m=0}^d \|X(i+m) - X(j+m)\|_2^2 \quad (17)$$

Let  $D_X^2$  be the  $N \times N$  matrix of all pairwise squared Euclidean distances between frames (possibly computed with the memory optimization in Section III-A), and let  $D_Y^2$  be the  $(N-d) \times (N-d)$  matrix of all pairwise distances between delay frames. Then Equation 17 implies that  $D_Y^2$  can be obtained from  $D_X^2$  via convolution with a “rect function”, or a vector of 1s of length  $d+1$ , over all diagonals in  $D_X^2$  (i.e. a moving average). This can be implemented in time  $O(N^2)$  with cumulative sums. Hence, regardless of how  $d$  is chosen, the computation and memory requirements for computing  $D_Y^2$  depend only on the number of frames in the video. Also,  $D_Y$  can simply be computed by taking the entry wise square root of  $D_Y^2$ , another  $O(N^2)$  computation. A similar scheme was used in [16] when comparing distances of 3D shape descriptors in videos of 3D meshes.

Figure 9 shows self-similarity matrices on embeddings of the pendulum video with no delay and with a delay approximately matching the period. The effect of a moving average along diagonals with delay eliminates the anti-diagonals caused by the video’s mirror symmetry.




Fig. 9: Self-similarity matrices  $D_{Y_{0,1}}$  and  $D_{Y_{28,1}}$  for a video of a the oscillating pendulum. Bright colors indicate far distances and dark colors indicate near distances. This example clearly shows how adding a delay embedding is like performing block averaging along all diagonals of the pairwise distance matrices, and it gets rid of the mirror symmetry.

Even for videos without mirror symmetries, such as a video of a running dog (Figure 10), introducing a delay brings the geometry into focus, as shown in Figure 11.

### C. Normalization

A few normalization steps are needed in order to enable fair comparisons between videos with different resolutions, or which have a different range in periodic motion either spatially or in intensity. First, we perform a “point-center and sphere normalize” vector normalization which was shown in [31] to have nice theoretical properties.

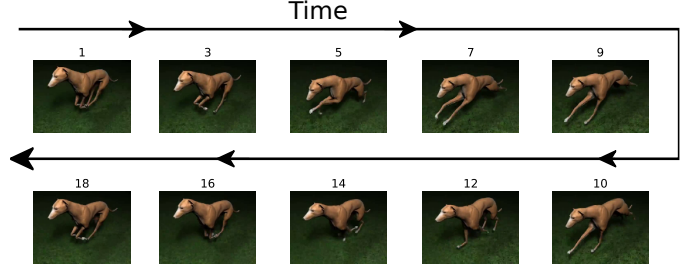


Fig. 10: An animation of a periodic video of a running dog, which, unlike an oscillating pendulum, does not have mirror symmetry in the second half of its period.




Fig. 11: Self-similarity matrices  $D_{Y_{0,1}}$  and  $D_{Y_{18,1}}$  for a video of a running dog. Even without the delay embedding ( $d = 0$ ), the video frames still form a topological loop. However, a delay embedding with  $d = 25$  cleans up the geometry and leads to a rounder loop, as seen in the resulting SSM.

That is,

$$\widetilde{SW}_{d,\tau}(t) = \frac{SW_{d,\tau}(t) - (SW_{d,\tau}(t)^T \mathbf{1}) \mathbf{1}}{\|SW_{d,\tau}(t) - (SW_{d,\tau}(t)^T \mathbf{1}) \mathbf{1}\|_2} \quad (18)$$

where  $\mathbf{1}$  is a  $WH(d+1) \times 1$  vector of all ones. In other words, one subtracts the mean of each component of each vector, and each vector is scaled so that it has unit norm (i.e. lives on the unit sphere in  $\mathbb{R}^{WH(d+1)}$ ). Subtracting the mean from each component will eliminate additive linear drift on top of the periodic motion, while scaling addresses resolution / magnitude differences. Note that we can still use the memory optimization in Section III-A, but we can no longer use the optimizations in Section III-B since each window is normalized independently.

Moreover, in order to mitigate nonlinear drift, we implement a simple pixel-wise convolution by the derivative of a Gaussian for each pixel in the original video before applying the delay embedding:

$$\hat{X}_i(t) = X_i(t) * -at \exp^{-t^2/(2\sigma^2)} \quad (19)$$

This is a pixel-wise bandpass filter which could be replaced with any other bandpass filter leveraging application specific knowledge of expected frequency bounds. This has the added advantage of reducing the number of harmonics, enabling a smaller embedding dimension  $d$ .

### D. Periodicity/Quasiperiodicity Scoring

Once the videos are normalized to the same scale, we can score periodicity and quasiperiodicity based on the ge-

ometry of sliding window embeddings. Let  $\text{dgm}_n$  be the  $n$ -dimensional persistence diagram for the Rips filtration on the sliding window embedding of a video, and define  $mp_i(\text{dgm}_n)$  as the  $i$ -th largest difference  $d - b$  for  $(b, d) \in \text{dgm}_n$ . In particular

$$mp_1(\text{dgm}_n) = \max\{d - b : (b, d) \in \text{dgm}_n\}$$

and  $mp_i(\text{dgm}_n) \geq mp_{i+1}(\text{dgm}_n)$ . We propose the following scores:

1) *Periodicity Score (PS)*

$$PS = \frac{1}{\sqrt{3}} mp_1(\text{dgm}_1) \quad (20)$$

Like [31], we exploit the fact that for the Rips filtration on  $S^1$ , the 1-dimensional persistence diagram has only one prominent birth-death pair with coordinates  $(0, \sqrt{3})$ . Since this is the limit shape of a normalized perfectly periodic sliding window video, the periodicity score is between 0 (not periodic) and 1 (perfectly periodic).

2) *Quasiperiodicity Score (QPS)*

$$QPS = \sqrt{\frac{mp_2(\text{dgm}_1) mp_1(\text{dgm}_2)}{3}} \quad (21)$$

This score is designed with the torus in mind. We score based on the second largest 1D persistence *times* the largest 2D persistence, since we want a shape that has two core circles and encloses a void to get a large score. Based on the Künneth theorem of homology, the 2-cycle (void) should die the moment the smallest 1-cycle dies.

3) *Modified Periodicity Score (MPS)*

$$MPS = \frac{1}{\sqrt{3}} (mp_1(\text{dgm}_1) - mp_2(\text{dgm}_1)) \quad (22)$$

We design a modified periodicity score which should be lower for quasiperiodic videos, than what the original periodicity score would yield.

Note that we use  $\mathbb{Z}_3$  field coefficients for all persistent homology computations since, as shown by [31], this works better for periodic signals with strong harmonics. Before we embark on experiments, let us explore the choice of two crucial parameters for the sliding window embedding: the delay  $\tau > 0$  and the dimension  $d \in \mathbb{N}$ . In practice we determine an equivalent pair of parameters: the dimension  $d$  and the *window size*  $d\tau$ .

### E. Dimension and Window Size

Takens' embedding theorem is one of the most fundamental results in the theory of dynamical systems [38]. In short, it contends that (under appropriate hypotheses) there exists an integer  $D$ , so that for all  $d \geq D$  and generic  $\tau > 0$  the sliding window embedding  $SW_{d,\tau}X$  reconstructs the state space of the underlying dynamics witnessed by the signal  $X$ . One common strategy for determining a minimal such  $D$  is the *false nearest-neighbors* scheme [21]. The idea is to keep track of the  $k$ -th nearest neighbors of each point in the delay embedding, and if they change as  $d$  is increased, then the prior

estimates for  $d$  were too low. This algorithm was used in recent work on video dynamics [42], for instance.

Even if we can estimate  $d$ , however, how does one choose the delay  $\tau$ ? As shown in [31], the sliding window embedding of a periodic signals is roundest (i.e. so that the periodicity score  $PS$  is maximized) when the window size,  $d\tau$ , satisfies the following relation:

$$d\tau = \frac{\pi k}{L} \left( \frac{d}{d+1} \right) \quad (23)$$

Here  $L$  is number of periods that the signal has in  $[0, 2\pi]$  and  $k \in \mathbb{N}$ . To verify this experimentally, we show in Figure 12 how the periodicity score  $PS$  changes as a function of window size for the pendulum video, and how the choice of window size from Equation 23 maximizes  $PS$ . To generate this figure we fixed a sufficiently large  $d$  and varied  $\tau$ . Let us now describe the general approach: Given a video we perform a period-length estimation step (see section III-F next), which results in a positive real number  $\ell$ . For a given  $d \in \mathbb{N}$  large enough we let  $\tau > 0$  be so that  $d\tau = \ell \cdot \frac{d}{d+1}$ .




Fig. 12: Varying the window size,  $d\tau$ , in a delay embedding of the synthetic pendulum video, which has a period length around 25 frames. Red dashed lines are drawn at the window lengths that would be expected to maximize roundness of the embedding for that period length based on theory in [31].

### F. Fundamental Frequency Estimation

Though Figure 12 suggests robustness to window size as long as the window is more than half of the period, we may not know what that is in practice. To automate window size choices, we do a coarse estimate using fundamental frequency estimation techniques on a 1D surrogate signal. To get a 1D signal, we extract the first coordinate of diffusion maps [4] using 10% nearest neighbors on the raw video frames (no delay) after taking a smoothed time derivative. Note that a similar diffusion-based method was also used in recent work by [46] to analyze the frequency spectrum of a video of an oscillating 2 pendulum + spring system in a quasiperiodic state. Once we have the diffusion time series, we then apply the normalized autocorrelation method of [25] to estimate the fundamental frequency. In particular, given a discrete signal  $x$  of length  $N$ , define the autocorrelation as

$$r_t(\tau) = \sum_{j=t}^{t+N-1-\tau} x_j x_{j+\tau} \quad (24)$$



However, as observed by [7], a more robust function for detecting periodicities is the *squared difference function*

$$d_t(\tau) = \sum_{j=t}^{t+N-1-\tau} (x_j - x_{j+\tau})^2 \quad (25)$$

which can be rewritten as  $d_t(\tau) = m_t(\tau) - 2r_t(\tau)$  where

$$m_t(\tau) = \sum_{j=t}^{t+N-1-\tau} (x_j^2 + x_{j+\tau}^2) \quad (26)$$

Finally, [25] suggest normalizing this function to the range  $[-1, 1]$  to control for window size and to have an interpretation akin to a Pearson correlation coefficient:

$$n_t(\tau) = 1 - \frac{m_t(\tau) - 2r_t(\tau)}{m_t(\tau)} = \frac{2r_t(\tau)}{m_t(\tau)} \quad (27)$$

The fundamental frequency is then the inverse period of the largest peak in  $n_t$  which is to the right of a zero crossing. The zero crossing condition helps prevent an offset of 0 from being the largest peak. Defining the normalized autocorrelation as in Equation 27 has the added advantage that the value of  $n_t(\tau)$  at the peak can be used as a way of scoring periodicity, which the authors call *clarity*. Values closer to 1 indicate more perfect periodicities. Note that this technique will sometimes pick integer multiples of the period. To mitigate this, we multiply  $n_t(\tau)$  by a slowly decaying envelope which is 1 for 0 lag and 0.9 for the maximum lag to emphasize smaller periods. Figure 13 shows the result of this algorithm on a periodic video, and Figure 14 shows the algorithm on an irregular video.




Fig. 13: Diffusion maps + normalized autocorrelation fundamental frequency estimation on a periodic vocal folds video (Section IV-C). The chosen period length is 32, as indicated by the red dot over the peak. This matches with the visually inspected period length.

#### IV. EXPERIMENTAL EVALUATION

Next we evaluate the effectiveness of the proposed (Modified) Periodicity and Quasiperiodicity scores on three different tasks. First, we provide estimates of accuracy for the binary classifications periodic/not-periodic or quasiperiodic/not-quasiperiodic in the presence of several noise models and noise levels. The results illustrate the robustness of our method. Second, we quantify the quality of periodicity rankings from





Fig. 14: Diffusion maps + normalized autocorrelation fundamental frequency estimation on a video of vocal folds with irregular oscillations (Section IV-C).

machine scores, as compared to those generated by human subjects. In a nutshell, and after comparing with several periodicity quantification algorithms, our approach is shown to be the most closely aligned with the perception of human subjects. Third, we demonstrate that our methodology can be used to automatically detect the physiological manifestations of certain speech pathologies (e.g., normal vs. biphonation), directly from high-speed videos of vibrating vocal folds.

##### A. Classification Under Varying Noise Levels/Models

As shown empirically in [8], a common source of noise in videos comes from camera shake (blur); this is captured by point spread functions resembling directed random walks [8, Figure 1] and the amount of blur (i.e. noise level) is controlled by the extent in pixels of the walk. Other sources are additive white Gaussian noise (awgn), controlled by the standard deviation of the Gaussian kernel, and MPEG bit errors quantified by the percentage of corrupted information. Figure 15 shows examples of the effects of these sources of noise.



(a) Original (b)  $20 \times 20$  blur (c) awgn  $\sigma = 2$  (d) 5% Bit Error

Fig. 15: The results of applying motion blur, additive white Gaussian noise, and MPEG bit corruption to a video frame.

For classification purposes we use three main recurrence classes. Three types of periodic videos (True periodic, TP): an oscillating pendulum, a bird flapping its wings, and an animation of a beating heart. Two types of quasiperiodic videos (True quasiperiodic, TQ): one showing two solid disks which oscillate sideways at non-commensurate rates, and the second

showing two stationary Gaussian pulses with amplitudes non-commensurately modulated by cosine functions. Two videos without significant recurrence (True non-recurrent, TN): a video of a car driving past a landscape, and a video of an explosion. Each one of these seven videos is then corrupted by the three noise models at three different noise levels (blur = 20, 40, 80, awgn  $\sigma = 1, 2, 3$ , bit error = 5, 10, 20%) as follows: given a particular video, a noise model and noise level, 600 instances are generated by sampling noise independently at random.

**Results:** We report in Table I the area under the Receiver Operating Characteristic (ROC) curve, or AUROC for short, for the classification task TP vs. TN (resp. TQ vs. TN) and binary classifier furnished by Periodicity (resp. Quasiperiodicity) Score.

For instance, for the Blur noise model with noise level of  $80 \times 80$  pixels, the AUROC from using the Periodicity Score to classify the 600 instances of the Heartbeat video as periodic, and the 600 instances of the Driving video as non-recurrent is 0.91. Similarly, for the MPEG bit corruption model with 5% of bit error, the AUROC from using the Quasiperiodicity score to classify the 600 instances of the Quasiperiodic Sideways Disks as quasiperiodic and the 600 instances of the Explosions video as non-recurrent is 0.92. To put these numbers in perspective, AUROC = 1 is associated with a perfect classifier and AUROC = 0.5 corresponds to classification by a random coin flip.

Overall, the type of noise that degrades performance the most across videos is the bit error, which makes sense, since this has the effect of randomly freezing, corrupting, or even deleting frames, which all interrupt periodicity. The blur noise also affects videos where the range of motion is small. The pendulum video, for instance, only moves over a range of 60 pixels at the most extreme end, so an  $80 \times 80$  pixel blur almost completely obscures the motion.

### B. Comparing Human and Machine Periodicity Rankings

Next we quantify the extent to which rankings obtained from our periodicity score (Equation 20), as well as three other methods, agree with how humans rank videos by periodicity. The starting point is a dataset of 20 different creative commons videos, each 5 seconds long at 30 frames per second. Some videos appear periodic, such as a person waving hands, a beating heart, and spinning carnival rides. Some of them appear nonperiodic, such as explosions, a traffic cam, and drone view of a boat sailing. And some of them are in between, such as the pendulum video with simulated camera shake.

It is known that humans are notoriously bad at generating globally consistent rankings of sets with more than 5 or 7 elements [27]. However, when it comes to binary comparisons of the type “should A be ranked higher than B?” few systems are as effective as human perception, specially for the identification of recurrent patterns in visual stimuli. We will leverage this to generate a globally consistent ranking of the 20 videos in our initial data set.

We use Amazon’s Mechanical Turk (AMT) [5] to present each pair of videos in the set of 20,  $\binom{20}{2} = 190$ , each to three different users, for a total of 570 pairwise rankings. 15

unique AMT workers contributed to our experiment <sup>2</sup>, using an interface as the one shown in Figure 16.

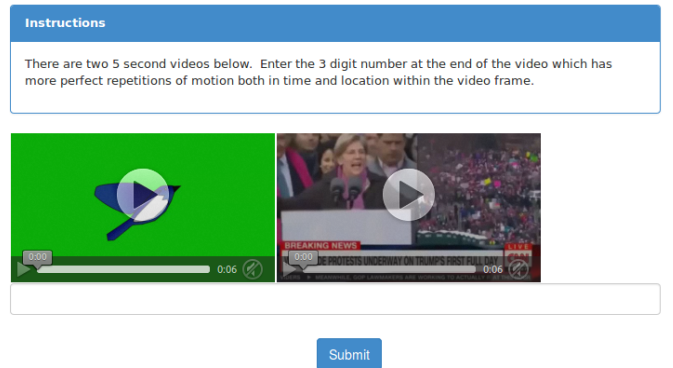


Fig. 16: The interface that humans are given on AMT for pairwise ranking videos by periodicity.

In order to aggregate this information into a global ranking which is as consistent as possible with the pairwise comparisons, we implement a technique known as *Hodge rank aggregation* [18]. Hodge rank aggregation finds the closest consistent ranking to a set of preferences, in a least squares sense. More precisely, given a set of objects  $X$ , and given a set of comparisons  $P \subset X \times X$ , we seek a scalar function  $s$  on all of the objects that minimizes the following sum

$$\sum_{(a,b) \in P} |v_{ab} - (s_b - s_a)|^2 \quad (28)$$

where  $v_{ab}$  is a real number which is positive if  $b$  is ranked higher than  $a$  and negative otherwise. Thus,  $s$  is a function whose discrete gradient best matches the set of preferences with respect to an  $L^2$  norm. Note that the preferences that we feed the algorithm are based on the pairwise rankings returned from AMT. If video  $b$  is greater than video  $a$ , then we assign  $v_{ab} = 1$ , or  $-1$  otherwise. Since we have 3 rankings for each video, we actually assign weights of  $+3$ ,  $+1$ ,  $-1$ , or  $-3$ . The  $\pm 3$  are if all rankings agree in one direction, and the  $\pm 1$  are if one of the rankings disagrees with the other two. Figure 17 shows a histogram of all of the weighted scores from users on AMT. They are mostly in agreement, though there are a few  $\pm 1$  scores.




Fig. 17: The histogram of scores that the workers on AMT gave to all pairwise videos.

As comparison to the human scores, we use three different classes of techniques for machine ranking of periodicity.

<sup>2</sup>We gratefully acknowledge the time and effort put in by anonymous people on the Internet to helps us to create an unbiased global ranking!

TABLE I: AUROC values for different levels of noise comparing periodic to non-periodic videos (driving - left subcell, explosions - right subcell) based on the periodic score (Equation 20). Also two synthetic quasiperiodic videos (sideways disk, modulated pulse) are compared to the same two non-periodic videos based on the quasiperiodic score (Equation 21).

	Blur 20		Blur 40		Blur 80		Awgn $\sigma = 1$		Awgn $\sigma = 2$		Awgn $\sigma = 3$		5% Bit Error		10% Bit Error		20% Bit Error	
Bird Flapping	1	1	1	1	1	1	1	1	1	1	1	1	1	0.97	1	0.92	0.75	0.79
Heartbeat	1	1	1	1	0.91	0.94	1	1	1	1	0.9	0.89	1	0.91	0.98	0.87	0.56	0.62
Pendulum	1	0.94	0.84	0.85	0.43	0.68	1	1	1	1	1	1	1	1	1	0.99	0.98	0.97
Quasiperiodic 1	1	0.85	1	0.83	0.75	0.39	1	1	1	1	1	1	1	0.92	0.99	0.93	0.82	0.82
Quasiperiodic 2	1	0.9	1	0.87	1	0.83	1	1	1	1	1	1	1	0.96	0.99	0.95	0.95	0.95

*Sliding Windows (SW)*: We sort the videos in decreasing order of Periodicity Score (Equation 20). We fix the window size at 20 frames and the embedding dimension at 20 frames (which is enough to capture 10 strong harmonics). We also apply a time derivative of width 10 to every frame.

*Cutler-Davis [6]*: This is one of the few techniques in the literature that can be readily used to quantify periodicity on a continuum. The authors use two different techniques to quantify periodicity from a self-similarity matrix (SSM) of video frames. The first is a frequency domain technique based on the peak of the average power spectral density over all columns (rows) of the SSM after linearly de-trending and applying a Hann window. To turn this into a continuous score, we report the ratio of the peak minus the mean over the standard deviation. This method will be referred to as *Frequency Score*.

As the authors warn, the frequency peak method has a high susceptibility to false positives. This motivated the design of a more robust technique in [6], which works by finding peaks in the 2D normalized autocorrelation of the Gaussian smoothed SSMs. For videos with mirror symmetry, the peaks will lie on a diamond lattice (Figure 18), while for videos without mirror symmetry, they will lie on a square lattice. After peak finding within neighborhoods, one simply searches over all possible lattices at all possible widths to find the best match with the peaks. Since each lattice is centered at the autocorrelation point  $(0, 0)$ , no translational checks are necessary.




Fig. 18: The method of [6] on the pendulum video. The peak PSD frequency is above two times the standard deviation above the mean (green line, center figure). The maxes (green points) on the 2D autocorrelation of the SSM lie on a diamond shape pattern.

To turn this into a continuous score, let  $E$  be the sum of Euclidean distances of the matched peaks in the autocorrelation image to the best fit lattice, let  $r_1$  be the proportion of lattice points that have been matched, and let  $r_2$  be the proportion

of peaks which have been matched to a lattice point. Then we give the final periodicity score as

$$CD_{score} = (1 + E/r_1)/(r_1 r_2)^3 \quad (29)$$

A lattice which fits the peaks perfectly ( $r_1 = 1$ ) with no error ( $E = 0$ ) and no false positive peaks ( $r_2 = 1$ ) will have a score of 1, and any video which fails to have a perfectly matched lattice will have a score greater than 1. Hence, we sort in increasing order of the score to get a ranking.

As we will show, this technique agrees the second best with humans after our Periodicity Score ranking. One of the main drawbacks is numerical stability of finding maxes in non-isolated critical points around nearly diagonal regions in square lattices, which will erroneously inflate the score. Also, the lattice searching only occurs over an integer grid, but there may be periods that aren't integer number of frames, so there will always be a nonzero  $E$  for such videos. By contrast, our sliding window scheme can work for any real valued period length.

*Diffusion Maps + Normalized Autocorrelation "Clarity"*: Finally, we apply the technique from Section III-F to get an autocorrelation function, and we report the value of the maximum peak of the normalized autocorrelation to the right of a zero crossing, referred to as "clarity" by [25]. Values closer to 1 indicate more perfect repetitions, so we sort in descending order of clarity to get a ranking.

Figure 19 shows an example of these three different techniques on a periodic video. There is a dot which rises above the diagonal in the persistence diagram, a lattice is found which nearly matches the critical points in the autocorrelation image, and autocorrelation function on diffusion maps has a nice peak.

By contrast, for a nonperiodic video (Figure 20), there is hardly any persistent homology, there is no well matching lattice, and the first diffusion coordinate has no apparent periodicities.

**Results:** Once we have the global human rankings and the global machine rankings, we can compare them using the *Kendall  $\tau$  score* [20]. Given a set of objects  $N$  objects  $X$  and two total orders  $>_1$  and  $>_2$ , where  $> (x_a, x_b) = 1$  if  $x_a > x_b$  and  $> (x_a, x_b) = -1$  if  $x_a < x_b$ , the Kendall  $\tau$  score is defined as

$$\tau = \frac{1}{N(N-1)/2} \sum_{i < j} (>_1(x_i, x_j))(>_2(x_i, x_j)) \quad (30)$$

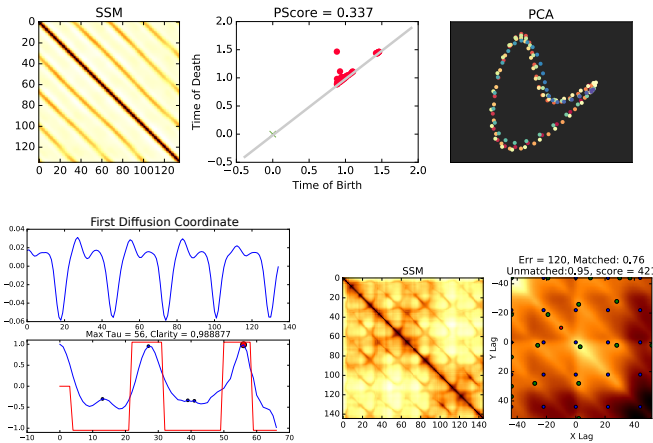


Fig. 19: An example of the SW score (top), the clarity score (bottom left), and the  $CD_{score}$  (bottom right, matched peaks in green and lattice in blue), on a periodic video of a man waving his arms from the KTH dataset ([36]).

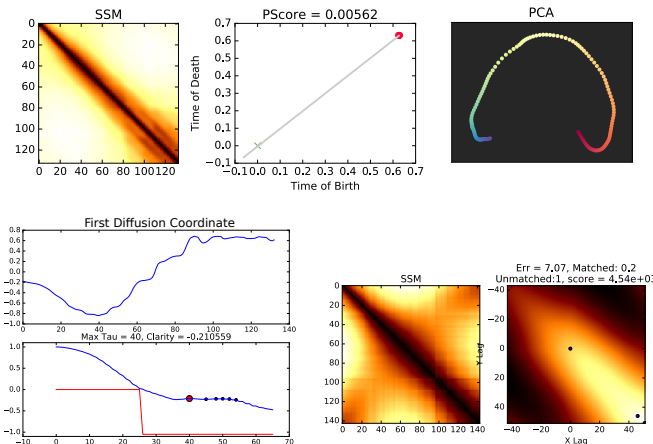


Fig. 20: An example of the SW score (top), the clarity score (bottom left), and the  $CD_{score}$  (bottom left, matched peaks in green and lattice in blue), on a video of an explosion, which is nonperiodic.

For two rankings which agree exactly, the Kendall  $\tau$  score will be 1. For two rankings which are exactly the reverse of each other, the Kendall  $\tau$  score will be -1. In this way, it is analogous to a Pearson correlation between rankings. Table II shows the Kendall  $\tau$  scores between all of the different machine rankings and the human rankings.

TABLE II: The Kendall  $\tau$  scores between all of the machine rankings and the Hodge aggregated human rankings.

	Human	SW	Freq	$CD_{score}$	Clarity
Human	1	<b>0.663</b>	-0.295	0.347	0.284
SW	0.663	1	-0.316	0.221	<b>0.516</b>
Freq	-0.295	-0.316	1	-0.0842	-0.189
$CD_{score}$	0.347	0.221	-0.0842	1	0.411
Clarity	0.284	0.516	-0.189	0.411	1

Our sliding window video methodology agrees with the human ranking more than any other pair of ranking types. The second most similar are the SW and the diffusion clarity,

which is noteworthy as they are both geometric techniques.

### C. Periodicity And Biphonation in High Speed Videos of Vocal Folds

In this final task we apply our methodology to a real world problem of interest in medicine. We show that our method can automatically detect certain types of voice pathologies from high-speed glottography, or high speed videos (4000 fps) of the left and right vocal folds in the human vocal tract [9], [44], [45]. In particular, we detect and differentiate quasiperiodicity from periodicity by using our geometric sliding window pipeline. Quasiperiodicity is a special case of what is referred to as “biphonation” in the biological context, where nonlinear phenomena cause a physical process to bifurcate into two different periodic modes, often during a transition to chaotic behavior [15]. The torus structure we sketched in Figure 8 has long been recognized in this context [14], [15], but we provide a novel way of quantifying it.

In contrast with our work, the existing literature usually employs an inherently *Lagrangian* approach, where different points on the left and right vocal folds are tracked, and coordinates of these points are analyzed as 1D time series (e.g. [13], [24], [28], [35], [26]). This is a natural approach, since those are the pixels where all of the important signal resides, and well-understood 1D signal processing technique can be used. However, edge detectors often require tuning, and they can suddenly fail when the vocal folds close [24]. In our technique, we give up the ability to localize the anomalies (left/right, anterior/posterior) since we are not tracking them, but in return we do virtually no preprocessing, and our technique is domain independent.

**Results:** We use a collection of 7 high-speed videos for this analysis, drawn from a variety of different sources [49], [26], [28], [13]. There are two videos which correspond to “normal” periodic vocal folds, three which correspond to biphonation [28], and two which correspond to irregular motion. We manually extracted 400 frames per video (100milliseconds) and autotuned the window size based on autocorrelation of 1D diffusion maps (Section III-F). We then chose an appropriate  $\tau$  and chose a time spacing so that each point cloud would have 600 points. As shown in Table III, our technique is able to differentiate between the four classes. We also show PCA and persistence diagrams for one example for each class. In Figure 21, we see what appears to be a loop in PCA, and one strong 1D persistent dot confirms this. In Figure 22, we see a prominent torus in the persistence diagram. In Figure 23, we don’t see any prominent structures in the persistence diagram, even though PCA looks like it could be a loop or a torus. Note, however, that PCA only preserves 13.7% of the variance in the signal, which is why high dimensional techniques are important to draw quantitative conclusions.

## V. DISCUSSION

We have shown in this work how applying sliding window embeddings to videos can be used to translate properties of the underlying dynamics into geometric features of the resulting



TABLE III: Results of our sliding window pipeline on videos of periodic vocal folds, biphonation, and irregularities. We give the max persistence periodicity score (PS), the modified periodicity score (MPS), the harmonic score (HS), and quasiperiodic score (QPS) presented in Section III-D. We also show the window size (Win) that the autocorrelation technique in Section III-F gives. We have bolded the top three MPS, HS, and QPS scores across all videos. The max modified periodic scores include the two periodic videos and one of the biphonation videos. The max quasiperiodic scores are all of the biphonation videos, which means the one with a high periodicity score could be ruled out of the periodicity category.

Video Name	Win	PS	MPS	QPS
Periodic 1 ([13])	16	0.816	<b>0.789</b>	0.011
Periodic 2 ([26], Figure 21)	32	0.601	<b>0.533</b>	0.009
Biphonation 1 ([28])	53	0.638	0.294	<b>0.292</b>
Biphonation 2 ([28])	42	0.703	<b>0.583</b>	<b>0.116</b>
Biphonation 3 ([28], Figure 22)	67	0.515	0.076	<b>0.426</b>
Mucus Perturbed Periodic ([49])	94	0.028	0.019	0.004
Irregular ([13], Figure 23)	232	0.18	0.097	0.04

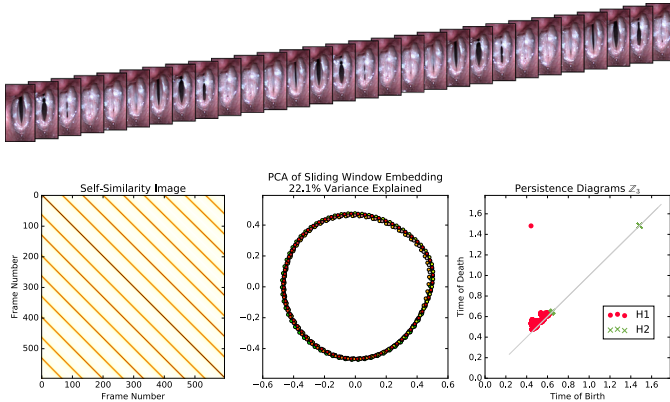


Fig. 21: Video frames and sliding window statistics on a video of vocal folds undergoing normal periodic vibrations [26]. One strong loop is visible in PCA and in the persistence diagrams

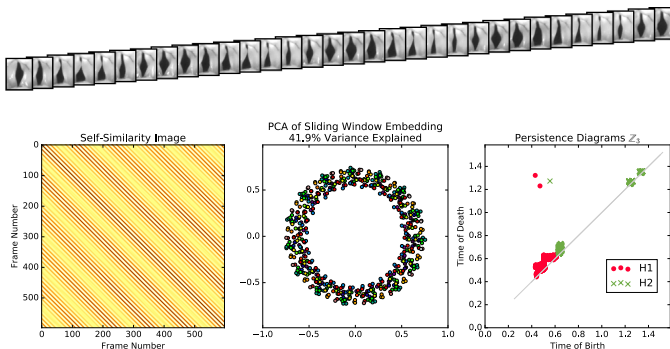


Fig. 22: Video frames and sliding window statistics on a video of vocal folds undergoing biphonation, courtesy of Juergen Neubauer [28]. PCA suggests a possible torus, and the persistence diagram indeed has the signature of a torus (two strong independent 1-cycles and one 2-cycle)

point cloud representation. Moreover, we also showed how topological/geometric tools such as persistence homology can




Fig. 23: Video frames and sliding window statistics of “irregular” vocal fold vibrations [13]. Though 2D PCA looks similar to Figure 22, no apparent 1D or 2D topological features are apparent in the high dimensional state space.

be leveraged to quantify the geometry of these embeddings. The pipeline was evaluated extensively showing robustness to several noise models, high quality in the produced periodicity rankings and applicability to the study of speech conditions from high-speed video data.

Moving forward, an interesting avenue related to medical applications is the difference between biphonation which occurs from quasiperiodic modes and biphonation which occurs from harmonic modes. [31] shows that  $\mathbb{Z}_3$  field coefficients can be used to indicate the presence of a strong harmonic, so we believe a geometric approach is possible. This could be used, for example, to differentiate between subharmonic anomalies and quasiperiodic transitions [44].

## ACKNOWLEDGMENTS

The authors would like to thank Juergen Neubauer, Dimitar Deliyiski, Robert Hillman, Alessandro de Alarcon, Dariush Mehta, and Stephanie Zacharias for providing videos of vocal folds. We also thank Matt Berger at ARFL for discussions about sliding window video efficiency. Christopher Tralie was partially supported by an NSF Graduate Fellowship NSF under grant DGF-1106401 and an NSF big data grant DKA-1447491. Jose Perea was partially supported by the NSF under grant DMS-1622301 and DARPA under grant HR0011-16-2-003.

## REFERENCES

- [1] Mark Allmen and Charles R Dyer. Cyclic motion detection using spatiotemporal surfaces and curves. In *Pattern Recognition, 1990. Proceedings., 10th International Conference on*, volume 1, pages 365–370. IEEE, 1990.
- [2] John Atanburi, Peter Cowling, John Murray, Belinda Colston, Paul Eady, Dave Hughes, Ian Nixon, and Patrick Dickinson. Analysis of bat wing beat frequency using fourier transform. In *International Conference on Computer Analysis of Images and Patterns*, pages 370–377. Springer, 2013.
- [3] Ulrich Bauer. Ripser: a lean C++ code for the computation of Vietoris-Rips persistence barcodes. <http://ripser.org>, 20152017.
- [4] Ronald R Coifman and Stéphane Lafon. Diffusion maps. *Applied and computational harmonic analysis*, 21(1):5–30, 2006.
- [5] Matthew JC Crump, John V McDonnell, and Todd M Gureckis. Evaluating amazon’s mechanical turk as a tool for experimental behavioral research. *PLoS one*, 8(3):e57410, 2013.
- [6] Ross Cutler and Larry S. Davis. Robust real-time periodic motion detection, analysis, and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):781–796, 2000.

- [7] Alain De Cheveigné and Hideki Kawahara. Yin, a fundamental frequency estimator for speech and music. *The Journal of the Acoustical Society of America*, 111(4):1917–1930, 2002.
- [8] Mauricio Delbracio and Guillermo Sapiro. Removing camera shake via weighted fourier burst accumulation. *IEEE Transactions on Image Processing*, 24(11):3293–3307, 2015.
- [9] Dimitar D Deliyski, Pencho P Petrushev, Heather Shaw Bonilha, Terri Treman Gerlach, Bonnie Martin-Harris, and Robert E Hillman. Clinical implementation of laryngeal high-speed videoendoscopy: challenges and evolution. *Folia Phoniatrica et Logopaedica*, 60(1):33–44, 2007.
- [10] Roman Goldenberg, Ron Kimmel, Ehud Rivlin, and Michael Rudzsky. Behavior classification by eigendecomposition of periodic motions. *Pattern Recognition*, 38(7):1033–1043, 2005.
- [11] Jerry P Gollub and Harry L Swinney. Onset of turbulence in a rotating fluid. *Physical Review Letters*, 35(14):927, 1975.
- [12] Allen Hatcher. *Algebraic topology*. University Press Ltd., 2002.
- [13] Christian T Herbst, Jakob Unger, Hanspeter Herzel, Jan G Švec, and Jörg Lohscheller. Phasegram analysis of vocal fold vibration documented with laryngeal high-speed video endoscopy. *Journal of Voice*, 30(6):771–e1, 2016.
- [14] Hanspeter Herzel, David Berry, Ingo R Titze, and Marwa Saleh. Analysis of vocal disorders with methods from nonlinear dynamics. *Journal of Speech, Language, and Hearing Research*, 37(5):1008–1019, 1994.
- [15] Hanspeter Herzel, Robert Reuter, and Richard A Katz. Biphonation in voice signals. In *AIP Conference Proceedings*, volume 375, pages 644–657. AIP, 1996.
- [16] Peng Huang, Adrian Hilton, and Jonathan Starck. Shape similarity for 3d video sequences of people. *International Journal of Computer Vision*, 89(2-3):362–381, 2010.
- [17] Shiyao Huang, Xianghua Ying, Jiangpeng Rong, Zeyu Shang, and Hongbin Zha. Camera calibration from periodic motion of a pedestrian. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3025–3033, 2016.
- [18] Xiaoye Jiang, Lek-Heng Lim, Yuan Yao, and Yinyu Ye. Statistical ranking and combinatorial hodge theory. *Mathematical Programming*, 127(1):203–244, 2011.
- [19] Holger Kantz and Thomas Schreiber. *Nonlinear time series analysis*, volume 7. Cambridge university press, 2004.
- [20] Maurice G Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938.
- [21] Matthew B Kennel, Reggie Brown, and Henry DI Abarbanel. Determining embedding dimension for phase-space reconstruction using a geometrical construction. *Physical review A*, 45(6):3403, 1992.
- [22] Orrawan Kumdee and Panrasee Ritthipravat. Repetitive motion detection for human behavior understanding from video images. In *Signal Processing and Information Technology (ISSPIT), 2015 IEEE International Symposium on*, pages 484–489. IEEE, 2015.
- [23] Ofir Levy and Lior Wolf. Live repetition counting. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3020–3028, 2015.
- [24] Jörg Lohscheller, Hikmet Toy, Frank Rosanowski, Ulrich Eysholdt, and Michael Döllinger. Clinically evaluated procedure for the reconstruction of vocal fold vibrations from endoscopic digital high-speed videos. *Medical image analysis*, 11(4):400–413, 2007.
- [25] Philip Mcleod and Geoff Wyvill. A smarter way to find pitch. In *In Proceedings of the International Computer Music Conference (ICMC05)*, pages 138–141, 2005.
- [26] Daryush D Mehta, Dimitar D Deliyski, Thomas F Quatieri, and Robert E Hillman. Automated measurement of vocal fold vibratory asymmetry from high-speed videoendoscopy recordings. *Journal of Speech, Language, and Hearing Research*, 54(1):47–54, 2011.
- [27] George A Miller. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological review*, 63(2):81, 1956.
- [28] Jürgen Neubauer, Patrick Mergell, Ulrich Eysholdt, and Hanspeter Herzel. Spatio-temporal analysis of irregular vocal fold oscillations: Biphonation due to desynchronization of spatial modes. *The Journal of the Acoustical Society of America*, 110(6):3179–3192, 2001.
- [29] Sourabh A Niyogi, Edward H Adelson, et al. Analyzing and recognizing walking figures in xyt. In *CVPR*, volume 94, pages 469–474, 1994.
- [30] Jose A Perea. Persistent homology of toroidal sliding window embeddings. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*, pages 6435–6439. IEEE, 2016.
- [31] Jose A Perea and John Harer. Sliding windows and persistence: An application of topological methods to signal analysis. *Foundations of Computational Mathematics*, 15(3):799–838, 2015.
- [32] Mark A Pinsky. *Introduction to Fourier analysis and wavelets*, volume 102. American Mathematical Soc., 2002.
- [33] Aaron M Plotnik and Stephen M Rock. Quantification of cyclic motion of marine animals from computer vision. In *OCEANS’02 MTS/IEEE*, volume 3, pages 1575–1581. IEEE, 2002.
- [34] Ramprasad Polana and Randal C Nelson. Detection and recognition of periodic, nonrigid motion. *International Journal of Computer Vision*, 23(3):261–282, 1997.
- [35] Qingjun Qiu, HK Schutte, Lide Gu, and Qilian Yu. An automatic method to quantify the vibration properties of human vocal folds via videokymography. *Folia Phoniatrica et Logopaedica*, 55(3):128–136, 2003.
- [36] Christian Schuldt, Ivan Laptev, and Barbara Caputo. Recognizing human actions: a local svm approach. In *Pattern Recognition, 2004. ICPR 2004. Proceedings of the 17th International Conference on*, volume 3, pages 32–36. IEEE, 2004.
- [37] Steven M Seitz and Charles R Dyer. View-invariant analysis of cyclic motion. *International Journal of Computer Vision*, 25(3):231–251, 1997.
- [38] Floris Takens. Detecting strange attractors in turbulence. In *Dynamical systems and turbulence, Warwick 1980*, pages 366–381. Springer, 1981.
- [39] Christopher Tralie. High-dimensional geometry of sliding window embeddings of periodic videos. In *LIPICs-Leibniz International Proceedings in Informatics*, volume 51. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016.
- [40] Matthew Turk and Alex Pentland. Eigenfaces for recognition. *Journal of cognitive neuroscience*, 3(1):71–86, 1991.
- [41] Mikael Vejdemo-Johansson, Florian T Pokorny, Primož Skraba, and Danica Kragic. Cohomological learning of periodic motion. *Applicable Algebra in Engineering, Communication and Computing*, 26(1-2):5–26, 2015.
- [42] V Venkataraman and P Turaga. Shape descriptions of nonlinear dynamical systems for video-based inference. *IEEE transactions on pattern analysis and machine intelligence*, 2016.
- [43] Ping Wang, Gregory D Abowd, and James M Rehg. Quasi-periodic event analysis for social game retrieval. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 112–119. IEEE, 2009.
- [44] Inka Wilden, Hanspeter Herzel, Gustav Peters, and Günter Tembrock. Subharmonics, biphonation, and deterministic chaos in mammal vocalization. *Bioacoustics*, 9(3):171–196, 1998.
- [45] Thomas Wittenberg, Manfred Moser, Monika Tigges, and Ulrich Eysholdt. Recording, processing, and analysis of digital high-speed sequences in glottography. *Machine vision and applications*, 8(6):399–404, 1995.
- [46] Or Yair, Ronen Talmon, Ronald R Coifman, and Ioannis G Kevrekidis. No equations, no parameters, no variables: data, and the reconstruction of normal forms by learning informed observation geometries. *arXiv preprint arXiv:1612.03195*, 2016.
- [47] Jing Yang, Hong Zhang, and Guohua Peng. Time-domain period detection in short-duration videos. *Signal, Image and Video Processing*, 10(4):695–702, 2016.
- [48] Guoshen Yu, Guillermo Sapiro, and Stéphane Mallat. Solving inverse problems with piecewise linear estimators: From gaussian mixture models to structured sparsity. *IEEE Transactions on Image Processing*, 21(5):2481–2499, 2012.
- [49] Stephanie RC Zacharias, Charles M Myer, Jareen Meinzen-Derr, Lisa Kelchner, Dimitar D Deliyski, and Alessandro de Alarcón. Comparison of videostroboscopy and high-speed videoendoscopy in evaluation of supraglottic phonation. *Annals of Otolaryngology, Rhinology & Laryngology*, page 0003489416656205, 2016.
- [50] A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, 2005.